White Paper on Multiple Independent Alerting Systems

James K. Kuchar

Department of Aeronautics and Astronautics

Massachusetts Institute of Technology

Cambridge, MA

July 13, 1998

This document provides an outline of the primary design and research issues relating to multiple, independent alerting systems. Following an overview of the basic system problem, the issues are presented notionally, and recommendations for further research are provided.

1. General Overview

Figure 1 shows a block-diagram model of a generic aerospace system. As shown, there are six main elements in the system architecture: Environment and Process; Sensors; Alerting Logic; Displays; Human Operator; and Controls. The task of the operator is to guide the process along a desired trajectory (e.g., fly to a destination) while making control inputs as needed due to changes in the environment (e.g., conflicting aircraft or weather). The states of the process and of the environment are monitored by the operator through a set of sensors and nominal information sources. These nominal information sources include system status cockpit displays, the view out the windscreen, communication by radio, etc. Under normal conditions, these nominal information sources are sufficient to allow the process to be controlled safely and efficiently. Note that the human operator in Fig. 1 may be a pilot (in which case the controlled process is the aircraft) or an air traffic controller (in which case the controlled process is the set of traffic in the controller's sector).

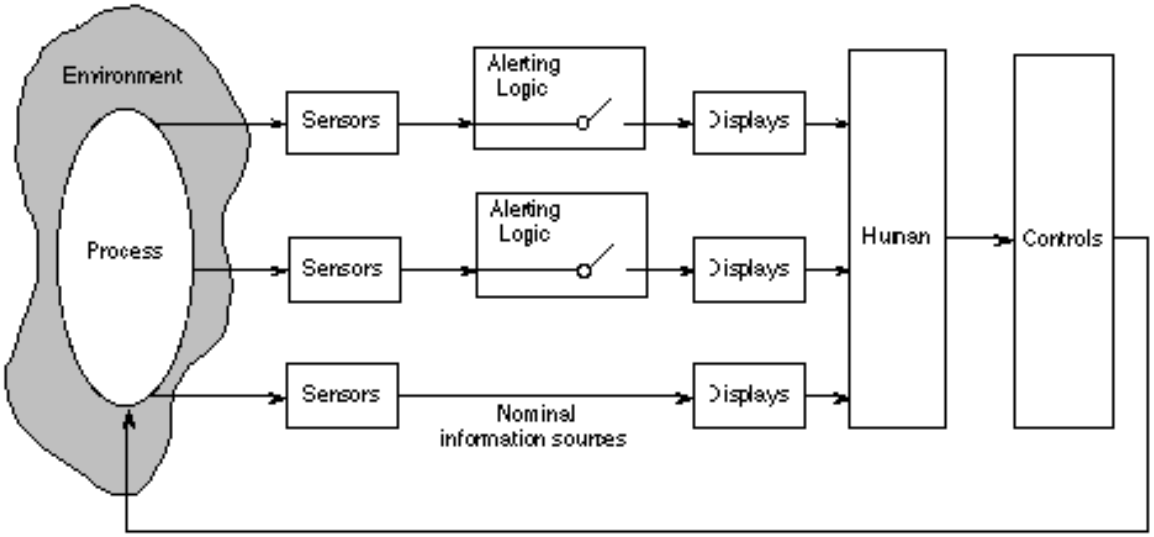

Figure 1: General Alerting System Architecture

Some events may occur that are not observable by the human (e.g., traffic conflict while in Instrument Conditions) or that may be observable but are not detected by the human (e.g., due to distraction with other tasks). Accordingly, alerting systems are typically introduced to monitor the process and environment in the background and alert the operator when certain criteria are met. These alerting systems generally use a separate suite of sensors and displays than are used with the nominal information sources. For example, TCAS uses transponder messages to estimate range and altitude whereas the pilot's nominal information source may be to observe traffic visually through the cockpit windscreen.

Until recently, alerting systems have been independent units, each targeted at a different hazard type (e.g., TCAS, GPWS, and engine fire warning system). Any overlap in responsibility between alerting systems was spread between different operators in the system. For example, conventional GPWS provides terrain alerts to flight crews, while Minimum Safe Altitude Warning Systems provide independent terrain alerts to ATC. In this way, additional redundancy can be introduced into the system.

The continued growth of automation on the flight deck and at Air Traffic Control facilities is beginning to result in overlap between independent alerting systems as viewed by a single operator. The most recent example is the Enhanced Ground Proximity Warning System (EGPWS), which uses independent sensor information to determine if a potential terrain threat exists. EGPWS operates entirely independently of conventional GPWS, and although EGPWS is designed to produce alerts before GPWS does, such behavior cannot be guaranteed.

Several proposed system enhancements will likewise involve overlap of responsibility. One example is a specialized collision alerting system that will operate during closely-spaced parallel approach, with conventional TCAS continuing to monitor other traffic in the area. A second area for overlap involves concepts for distributed conflict detection on the ground and in the cockpit, such as TCAS, CTAS, or future conflict probe and self-separation concepts.

Alerting systems generally perform four functions: hazard detection; attention-getting; display of resolution status and commands; and resolution guidance. Hazard detection and attention-getting involve developing decision thresholds upon which the decision is made as to whether a condition is hazardous. Issues relating to hazard detection and attention-getting are covered in Section 3. Section 4 focuses on resolution commands and guidance information when using two or more independent alerting systems. This includes information designed to aid the operator in selecting and following a course of action to safely and efficiently resolve the hazard situation.

2. Scope

This paper is focused on cases in which there are two or more alerting systems that have overlapping responsibility for a certain hazard category. It is assumed that the two systems use different state information or different alerting logic (e.g., TCAS TA and RA are not considered different systems, whereas EGPWS and GPWS are). Issues of prioritization between systems that monitor and alert for *different* hazard categories (e.g., engine fire vs. terrain proximity) are not included here.

A number of issues remain to be resolved regarding general alerting system design and

evaluation (e.g., tradeoffs between nuisance alarms and safety). However, this paper focuses only on that subset of issues specific to multiple, overlapping alerting systems. Issues that apply to single alerting systems are not included.

Additionally, although a range of human factors issues have been raised regarding alerting system design [1-7], the focus here is on the system architecture and algorithmic issues, not on the human factors issues.

3. Operational Modes for Alerting

The hazard detection process involves monitoring the states of the process and environment, and determining that a hazardous condition exists when the states exceed a predefined alert threshold. When this occurs, an attention-getting signal is transmitted so that the human operator is aware of the situation. In general, there are three functional modes in which two (or more) alerting systems may perform the hazard detection and attention-getting process, outlined below.

3.1 Simultaneous Operation

Both alerting systems may be operated simultaneously while monitoring the same hazard. An example of this mode of operation is GPWS and EGPWS. EGPWS provides an additional set of sensors, algorithms, and displays and is an independent enhancement to conventional GPWS (whose functionality is not changed). Each system continues to operate regardless of the status of the other system.

One advantage of this method is that the two systems provide additional redundancy in threat detection. Should one approach to alert detection not adequately model a particular hazard encounter, the second system may be able to detect and manage the threat. On the other hand, one may question why two systems are required to perform a task that ideally should be integrated into a single system. However, certification and training processes may be simplified when introducing an independent system enhancement rather than modifying and enhancing an existing system.

Potential issues in the simultaneous operation mode are generally related to mismatches in the timing or conditions under which alerts occur from each system. If both systems indicate that a hazardous state exists, there is no problem unless the system provides inconsistent resolution information. However, if one system indicates that there is a hazard while the other system does not, then the operator must have some basis with which to decide which system to believe. This basis may be conveyed through proper training in which the criteria and behavior of the two alerting systems is described. In some cases such a mismatch may be intentional in order to provide a progression of hazard levels. For example, EGPWS is designed to issue an alert before GPWS, in order to provide the pilot advance warning of a terrain threat. The pilots are taught this behavior and understand that the two systems are independent. Thus, an EGPWS alert without a GPWS alert is not abnormal, and if action is not taken, a GPWS alert is likely to occur shortly. However, threatening terrain could appear on an EGPWS display without any supporting GPWS alerts; the pilot is then left to determine which system to trust.

A designed progression of hazard levels between alerting systems could also fail to occur as intended. For example, it is possible that GPWS could alert before EGPWS (e.g., due to a terrain

database error), in which case the pilot may not fully trust the GPWS information.

The progression of alerts between two systems can be illustrated using a state-space diagram such as that shown in Fig. 2. Two systems, A and B, are each independently monitoring some potential hazard. System A uses one set of sensors to measure a state variable, x (e.g., time to impact). If x exceeds some threshold, an alert is issued from system A. Similarly, system B uses a different set of sensors to measure a different state variable, y (e.g., current range), and an alert is issued if y exceeds some other threshold. Depending on the relative rate at which x and y change, either system may alert first.

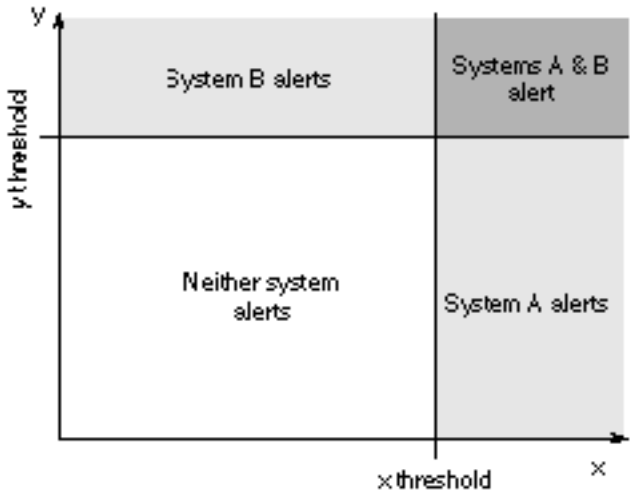

Figure 2: Multiple Alerting System Threshold Regions

Of particular concern would be a situation in which one system is increasing its hazard level at the same time that the second system is downgrading its hazard level. The pilot would then have to determine which trend was correct, and then act accordingly.

3.2 Complete inhibition

This mode of operation could occur, for example, in a system concept in which TCAS is entirely inhibited once a parallel approach begins. A second, independent system then takes over responsibility for collision alerting and resolution. This method would likely be used in a case in which specialized alerting systems are required in different phases of flight. For example, TCAS, though effective in managing enroute and terminal area traffic threats, is not optimized for closely-spaced parallel approaches and would therefore need to be inhibited to prevent nuisance alerts. The transition between systems may be designed to occur automatically or it may be designed to occur only upon command or selection of certain modes by the operator.

The advantage of this method is that the potential for simultaneous alerts from two systems is reduced. Additionally, because specialized alerting systems are used only in the phases of flight for which they were designed, more effective performance may be attained (in terms of nuisance alerts and safety level) than is possible with two more generic systems.

It is important, however, that the transition between the two systems is unambiguous and does

not produce discrepancies in threat level. This may require that the operator is informed as to which system is providing protection – potentially resulting in additional, undesirable displays that increase operator workload. Second, when a transition between systems occurs, it is necessary to ensure that the threat level remains constant or changes in a consistent manner. For example, a case in which one system produces an alert shortly before a transition to a second system (which no longer considers the hazard a threat) would be inappropriate.

In order to implement this mode of operation, specific criteria must be developed by which the transitions between systems are managed. This includes considerations of timing, hysteresis effects, and any necessary displays to enable the operator to understand the status of the alerting systems. Whether the operator needs to be informed when a transition occurs may depend on the type of alerting information that each system presents (discussed in more detail in Section 4).

3.3 Partitioned operation

In this mode of operation, both systems continue to operate, though each specific hazard in the envrionment is tracked by only one of the two systems. This method has been proposed for closely-spaced parallel approach. In the concept, a specialized alerting system would monitor only the relevant parallel traffic, while TCAS would continue to monitor other traffic in the area. Thus, depending on the source of a traffic conflict (parallel traffic or other traffic in the area) either a TCAS alert or a parallel approach alert could be generated. This is in contrast to the complete inhibition mode in which *all* hazards are tracked by one system.

This method has similar advantages to the mode in Section 3.2, though the partitioned mode allows for additional specialization of the alerting systems. Rather than be specialized for a given phase of flight, each system can be specialized for a particular type of hazard. This may allow better hazard management and further reduce unnecessary alerts and increase safety levels.

In parallel, similar disadvantages and design challenges apply to the partitioned mode as apply to the mode in Section 3.2. Requirements on consistent transitions between alerting systems may be more difficult in the partitioned mode, however, because rather than a single transition occurring, there may be continuous transitions as hazards are monitored first by one system and then the other. This places additional burdens on the design to provide hysteresis and to keep the operator aware of which hazard is being monitored by which system. Additionally, there may be conflicing resolution information between the two systems should each system alert against different threats. This issue is discussed in more detail in Section 4.

4. Operational Modes for Resolution

Resolution status, command, and guidance information is typically presented along with or shortly after attention-getting information. Status information is designed to inform the operator that there is a hazardous state of the process and environment (e.g., "Traffic" or "Terrain"). Command and guidance information is generally reserved for time-critical warnings in which the operator must act with a minimum of delay. Example commands include "Pull Up" or "Climb Climb" alerts; example guidance information is the presentation of TCAS RA pitch bars on the Primary Flight Display so that the pilot can monitor aircraft pitch and the required pitch angle to resolve the conflict.

There are three primary modes by which resolution information can be managed with two or more independent alerting systems, as outlined below.

4.1 Identical resolution information

In this mode, both alerting systems produce the same form of resolution information to the operator. Although the basis for generating the resolution information may vary between the alerting systems, the information presented to the operator does not distiguish which alerting system was the source. The benefit of this mode is that the underlying complexities of having two alerting systems is transparent to the operator. If there is no explicit need to differentiate between the two systems, then it is operationally more simple to instead generate a single type of resolution information, regardless of source. It may, however, in some applications be necessary to provide a distinction between systems, in which case this mode of operation cannot be used.

4.2 One system presents status, one system presents guidance / command

In this mode, one of the alerting systems is designed to provide status or situation awareness information, while a second alerting system is responsible for guidance and command information. One operational example using this mode is the EGPWS / GPWS combination, in which EGPWS does not provide resolution commands. The benefit of this mode of operation is that the two systems are more clearly separable, and may more easily be certified as system enhancements. It does, however, require additional operator training, as new outputs are now produced which the operator must interpet.

4.3 Both systems resolve using independent methods

Under this concept, the two alerting systems provide independent resolution information to the operator. Some form of prioritization is necessary in the event that both systems issue alerts at the same time. Potential problems include inconsistent or contradictory resolution information (e.g., one system commands a climb while the second system commands a descent).

It is also possible that the resolution information involves different types of action (e.g., one system assumes a standard procedure will be performed while the second system provides resolution guidance information). Again, care must be taken to ensure that the procedural maneuver is consistent with other guidance information. For example, if on system is designed assuming a climbing-turn procedural maneuver and a TCAS climb RA occurs against other traffic, the pilot must have a means of determining whether the procedure or RA has priority – ideally these should be supportive of each other, but this may not always be the case. It may be the case that by performing the climbing-turn, adequate vertical separation is not guaranteed by TCAS since TCAS assumes the maneuver is performed without the turning component.

Note that the Status/Guidance mode outlined in Section 4.2 is essentially a special case of the independent mode discussed here. The Status/Guidance mode, however, deserves specific mention because it clearly separates the type of information provided to the operator. In the independent method described here, more overlap among status or guidance information is assumed to exist.

5. Combinations of Modes of Operation

Given that there are three modes in which alerting may be performed and three modes in which resolution information may be presented, there is an array of possible system architectures. This array is shown in Table 1. Also shown in the table are examples of the EGPWS/GPWS and TCAS/AILS architectures.

Table 1: System Architecture Matrix

Resolution Mode	Alerting Mode: Simultaneous	Inhibited	Partitioned
Identical			
Status / Guidance	EGPWS/GPWS		
Independent			TCAS/AILS

More complex system architectures are certainly possible. For example, some systems may transition between cells in Table 1 (e.g., operating in Simultaneous mode in some phases of flight and in Inhibited or Partitioned mode in others). These mode transitions themselves have corresponding design issues relating to the methods by which the transitions are effected and displayed.

6. Areas for Additional Research

Based on the issues outlined above and the fact that there is little operational experience with these types of systems, the following areas for further research are recommended:

6.1 Transitions Between Modes in Inhibiting or Partitioned Systems

The potential problems presented here involving the management of the transition between systems are notional; little supporting data exists upon which to make design decisions. Studies are needed that better define the criticality of these issues. This includes examining appropriate criteria for determining when and how to transition between systems, when to initiate or disable inhibits, and to determine the need of an operator to know which system is operating and which hazards are being monitored by each system. Although many of these issues are application-specific, there are areas for underlying fundamental research that would provide needed insight.

6.2 Management of Independent Resolution Information

Research into resolution coordination between independent systems is also needed. This includes determining appropriate roles for procedural and guidance resolution information and ways in which these methods can be coordinated. Potential areas for conflict between resolution commands as well as guidelines for prioritization need to be determined.

6.3 Appropriate System Architectures

The matrix shown in Table 1 needs to be explored to determine which alerting and resolution mode is most effective for different applications. This includes identifying the criteria upon which the decision to use a given mode of alerting or resolution should be based. Additionally,

the design and research issues specific to each cell in the matrix need to be determined in more detail than is possible in this paper.

7. References

1. Veitengruber, Boucek, and Smith, "Aircraft Alerting Systems Criteria Study, Vol. 1: Collation and Analysis of Aircraft Alerting Systems Data", FAA-RD-76-222, May, 1977.

2. Boucek, Erickson, Berson, Hanson, Leffler, and Po-Chedley, "Aircraft Alerting Systems Standardization Study Phase I Final Report", FAA-RD-80-68, February, 1980.

3. Boucek, Po-Chedley, Berson, Hanson, Leffler, and White, "Aircraft Alerting Systems Standardization Study, Vol 1: Candidate System Validation and Time-Critical Display Evaluation", DOT/FAA/RD-81/38/I, January, 1981

4. Berson, Po-Chedley, Boucek, Hanson, Leffler, and Wasson, "Aircraft Alerting Systems Standardization Study, Vol 2: Aircraft Alerting System Design Guidelines", DOT/FAA/RD-81/38/II, January, 1981

5. Boucek, Berson, Po-Chedley, and Hendrickson, "Aircraft Alerting Systems Standardization Study", 4th AIAA/IEEE Digital Avionics Systems Conference, St. Louis, MO, November 17-19, 1981

6. Hanson, Howison, Chikos, and Berson, "Aircraft Alerting Systems Standardization Study, Phase IV: Accident Implications on System Design", DOT/FAA/RD-82/26, June, 1982.

7. Pritchett, A. R., "Pilot Non-Conformance to Alerting System Commands During Closely Spaced Parallel Approaches", Ph.D. Thesis, Department of Aeronautics and Astronautics, MIT, Cambridge, MA, December, 1996